

A Time-Efficient Quantum Walk for 3-Distinctness Using Nested Updates*

Andrew M. Childs^{†1,3}, Stacey Jeffery^{‡2,3}, Robin Kothari^{§2,3}, and Frédéric Magniez^{¶4}

¹Department of Combinatorics & Optimization, University of Waterloo, Canada

²David R. Cheriton School of Computer Science, University of Waterloo, Canada

³Institute for Quantum Computing, University of Waterloo, Canada

⁴CNRS, LIAFA, Univ Paris Diderot, Sorbonne Paris-Cité, France

Abstract

We present an extension to the quantum walk search framework that facilitates quantum walks with nested updates. We apply it to give a quantum walk algorithm for 3-Distinctness with query complexity $\tilde{O}(n^{5/7})$, matching the best known upper bound (obtained via learning graphs) up to log factors. Furthermore, our algorithm has time complexity $\tilde{O}(n^{5/7})$, improving the previous $\tilde{O}(n^{3/4})$.

1 Introduction

Element Distinctness is a basic computational problem. Given a sequence $\chi = \chi_1, \dots, \chi_n$ of n integers, the task is to decide if those elements are pairwise distinct. This problem is closely related to Collision, a fundamental problem in cryptanalysis. Given a 2-to-1 function $f : [n] \rightarrow [n]$, the aim is to find $a \neq b$ such that $f(a) = f(b)$. One of the best (classical and quantum) algorithms is to run Element Distinctness on f restricted to a random subset of size \sqrt{n} .

In the quantum setting, Element Distinctness has received a lot of attention. The first non-trivial algorithm used $\tilde{O}(n^{3/4})$ time [BDH⁺05]. The optimal $\tilde{O}(n^{2/3})$ algorithm is due to Ambainis [Amb04], who introduced an approach based on quantum walks that has become a major tool for quantum query algorithms. The optimality of this algorithm follows from a query lower bound for Collision [AS04]. In the query model, access to the input χ is provided by an oracle whose answer to query $i \in [n]$ is χ_i . This model is the quantum analog of classical decision tree complexity: the only resource measured is the number of queries to the input.

Quantum query complexity has been a very successful model for studying the power of quantum computation. In particular, quantum query complexity has been exactly characterized in terms of

*Support for this work was provided by NSERC, the Ontario Ministry of Research and Innovation, the US ARO, the French ANR Blanc project ANR-12-BS02-005 (RDAM), and the European Commission IST STREP project 25596 (QCS).

[†]amchilds@uwaterloo.ca

[‡]sjeffery@uwaterloo.ca

[§]rkothari@cs.uwaterloo.ca

[¶]frederic.magniez@univ-paris-diderot.fr

a semidefinite program, the general adversary bound [Rei11, LMR⁺11]. To design quantum query algorithms, it suffices to exhibit a solution to this semidefinite program. However, this turns out to be difficult in general, as the minimization form of the general adversary bound has exponentially many constraints. Belovs [Bel12b] recently introduced the model of learning graphs, which can be viewed as the minimization form of the general adversary bound with additional structure imposed on the form of the solution. This additional structure makes learning graphs much easier to reason about. The learning graph model has already been used to improve the query complexity of many graph problems [Bel12b, LMS11, LMS13] as well as k -Distinctness [Bel12a].

One shortcoming of learning graphs is that these upper bounds do not lead explicitly to efficient algorithms in terms of time complexity. Although the study of query complexity is interesting on its own, it is relevant in practice only when a query lower bound is close to the best known time complexity.

Recently, [JKM13] reproduced several known learning graph upper bounds via explicit algorithms in an extension of the quantum walk search framework of [MNRS11]. This work produced a new quantum algorithmic tool, quantum walks with nested checking. Algorithms constructed in the framework of [JKM13] can be interpreted as quantum analogs of randomized algorithms, so they are simple to design and analyze for any notion of cost, including time as well as query complexity. This framework has interpreted all known learning graphs as quantum walks, except the very recent *adaptive learning graphs* for k -Distinctness [Bel12a].

In k -Distinctness, the problem is to decide if there are k copies of the same element in the input, with $k = 2$ being Element Distinctness. The best lower bound for k -Distinctness is the Element Distinctness lower bound $\Omega(n^{2/3})$, whereas the best query upper bound is $O(n^{1-2^{k-2}/(2^k-1)}) = o(n^{3/4})$ [Bel12a], achieved using learning graphs, improving the previous bound of $O(n^{k/(k+1)})$ [Amb04]. However, the best known time complexity remained $\tilde{O}(n^{k/(k+1)})$. We improve this upper bound for the case when $k = 3$.

Our algorithm for 3-Distinctness is conceptually simple: we walk on sets of 2-collisions and look for a set containing a 2-collision that is part of a 3-collision. We check if a set has this property by searching for an index that evaluates to the same value as one of the 2-collisions in the set. However, to move to a new set of 2-collisions, we need to use a quantum walk subroutine for finding 2-collisions as part of our update step. This simple idea is surprisingly difficult to implement and leads us to develop a new extension of the quantum walk search framework.

Given a Markov chain P with spectral gap δ and success probability ε in its stationary distribution, one can construct a quantum search algorithm with cost $S + \frac{1}{\sqrt{\varepsilon}}(\frac{1}{\sqrt{\delta}}U + C)$ [MNRS11], where S , U and C are respectively the setup, update, and checking costs of the quantum analog of P . Using a quantum walk algorithm with costs $S', U', C', \varepsilon', \delta'$ (as in [MNRS11]) as a checking subroutine straightforwardly gives complexity $S + \frac{1}{\sqrt{\varepsilon}}(\frac{1}{\sqrt{\delta}}U + S' + \frac{1}{\sqrt{\varepsilon'}}(\frac{1}{\sqrt{\delta'}}U' + C'))$. Using nested checking [JKM13], the cost can be reduced to $S + S' + \frac{1}{\sqrt{\varepsilon}}(\frac{1}{\sqrt{\delta}}U + \frac{1}{\sqrt{\varepsilon'}}(\frac{1}{\sqrt{\delta'}}U' + C'))$.

It is natural to ask if a quantum walk subroutine can be used for the update step in a similar manner to obtain cost $S + S' + \frac{1}{\sqrt{\varepsilon}}(\frac{1}{\sqrt{\delta}}\frac{1}{\sqrt{\varepsilon'}}(\frac{1}{\sqrt{\delta'}}U' + C') + C)$. In most applications, the underlying walk is independent of the input, so the update operation is simple, but for some applications a more complex update may be useful (as in [CK11], where Grover search is used for the update). In Section 2.3, we describe an example showing that it is not even clear how to use a nested quantum walk for the update with the seemingly trivial cost $S + \frac{1}{\sqrt{\varepsilon}}(\frac{1}{\sqrt{\delta}}(S' + \frac{1}{\sqrt{\varepsilon'}}(\frac{1}{\sqrt{\delta'}}U' + C')) + C)$. Nevertheless, despite the difficulties that arise in implementing nested updates, we show in Section 3.2 how to achieve the more desirable cost expression in certain cases, and a similar one in general.

To accomplish this, we extend the quantum walk search framework by introducing the concept of *coin-dependent data*. This allows us to implement nested updates, with a quantum walk subroutine to carrying out the update procedure. Superficially, our modification appears small. Indeed, the proof of the complexity of our framework is nearly the same as that of [MNRS11]. However, there are some subtle differences in the implementation of the walk.

As in [JKM13], this concept is simple yet powerful. We demonstrate this by constructing a quantum walk version of the learning graph for 3-Distinctness with matching query complexity (up to poly-logarithmic factors). Because quantum walks are easy to analyze, the time complexity, which matches the query complexity, follows easily, answering an open problem of [Bel12a].

Independently, Belovs [Bel13] also recently obtained a time-efficient implementation of his learning graph for 3-Distinctness. His approach also uses quantum walks, but beyond this similarity, the algorithm appears quite different. In particular, it is based on another framework of search via quantum walk due to Szegedy [Sze04, MNRS12], whereas our approach uses a new extension of the quantum walk search framework of [MNRS11].

2 Preliminaries and Motivation

2.1 Quantum Walks

Consider a reversible, ergodic Markov chain P on a connected, undirected graph $G = (X, E)$ with spectral gap $\delta > 0$ and stationary distribution π . Let $M \subseteq X$ be a set of marked vertices. Our goal is to detect whether $M = \emptyset$ or $\Pr_{x \sim \pi}(x \in M) \geq \varepsilon$, for some given $\varepsilon > 0$. Consider the following randomized algorithm that finds a marked element with bounded error.

1. Sample x from π
2. Repeat for $\Theta(1/\varepsilon)$ steps
 - (a) If the current vertex x is marked, then stop and output x
 - (b) Otherwise, simulate $\Theta(1/\delta)$ steps of P starting with x
3. If the algorithm has not terminated, output ‘no marked element’

This algorithm has been quantized by [MNRS11] leading to efficient quantum query algorithms. Since each step has to be unitary and therefore reversible, we have to implement the walk carefully. The quantization considers P as a walk on edges of E . We write $(x, y) \in \vec{E}$ when we consider an edge $\{x, y\} \in E$ with orientation (x, y) . The notation (x, y) intuitively means that the current vertex of the walk is x and the *coin*, indicating the next move, is y . Swapping x and y changes the current vertex to y ; then the coin becomes x .

The quantum algorithm may carry some data structure while walking on G ; we formalize this as follows. Let 0 be a state outside X . Define $D : X \cup \{0\} \rightarrow \mathcal{D}$ for some Hilbert space \mathcal{D} , with $|D(0)\rangle = |0\rangle$. We define costs associated with the main steps of the algorithm. By cost we mean any measure of complexity such as query, time or space.

Setup cost: Let S be the cost of constructing

$$|\pi\rangle = \sum_{x \in X} \sqrt{\pi(x)} |x\rangle |D(x)\rangle \sum_{y \in X} \sqrt{P(x, y)} |y\rangle |D(y)\rangle.$$

Update cost: Let U be the cost of the LOCAL DIFFUSION operation, which is controlled on the first two registers¹ and acts as

$$|x\rangle |D(x)\rangle |0\rangle |D(0)\rangle \mapsto |x\rangle |D(x)\rangle \sum_{y \in X} \sqrt{P(x,y)} |y\rangle |D(y)\rangle.$$

Checking cost: Let C be the cost of the reflection

$$|x\rangle |D(x)\rangle \mapsto \begin{cases} -|x\rangle |D(x)\rangle & \text{if } x \in M \\ |x\rangle |D(x)\rangle & \text{otherwise.} \end{cases}$$

Theorem 2.1 ([MNRS11]). *Let P be a reversible, ergodic Markov chain on $G = (X, E)$ with spectral gap $\delta > 0$. Let $M \subseteq X$ be such that $\Pr_{x \sim \pi}(x \in M) \geq \varepsilon$, for some $\varepsilon > 0$, whenever $M \neq \emptyset$. Then there is a quantum algorithm that finds an element of M , if $M \neq \emptyset$, with bounded error and with cost*

$$O\left(S + \frac{1}{\sqrt{\varepsilon}} \left(\frac{1}{\sqrt{\delta}} U + C\right)\right).$$

Furthermore, we can approximately map $|\pi\rangle$ to $|\pi(M)\rangle$, the normalized projection of $|\pi\rangle$ onto $\text{span}\{|x\rangle |D(x)\rangle |y\rangle |D(y)\rangle : x \in M, y \in X\}$, in cost $\frac{1}{\sqrt{\varepsilon}}(\frac{1}{\sqrt{\delta}} U + C)$.

2.2 3-Distinctness

We suppose that the input is a sequence $\chi = \chi_1, \dots, \chi_n$ of integers from $[q] := \{1, \dots, q\}$. We model the input as an oracle whose answer to query $i \in [n]$ is χ_i .

We make the simplifying assumptions that there is at most one 3-collision and that the number of 2-collisions is in $\Theta(n)$. The first assumption is justified in [Amb04, Section 5]. To justify the second assumption, note that given an input $\chi \in [q]^n$, we can construct $\chi' \in [q+n]^{3n}$ with the same 3-collisions as χ , and $\Omega(n)$ 2-collisions, by defining $\chi'_i = \chi_i$ for $i \in [n]$ and $\chi'_i = \chi'_{i+n} = q+i$ for $i \in \{n+1, \dots, 2n\}$. Note that any two 2-collisions not both part of the 3-collision are disjoint.

A common simplifying technique is to randomly partition the space $[n]$ and assume that the solution respects the partition in some sense. Here we partition the space into three disjoint sets of equal size, A_1 , A_2 and A_3 , and assume that if there is a 3-collision $\{i, j, k\}$, then we have $i \in A_1$, $j \in A_2$ and $k \in A_3$. This assumption holds with constant probability, so we need only repeat the algorithm $O(1)$ times with independent choices of the tripartition to find any 3-collision with high probability. Thus, we assume we have such a partition.

2.3 Motivating Example

Quantum Walk for Element Distinctness. In the groundbreaking work of Ambainis [Amb04], which inspired a series of quantum walk frameworks [Sze04, MNRS11, JKM13] leading up to this work, a quantum walk for solving Element Distinctness was presented. This walk takes place on a Johnson graph, $J(n, r)$, whose vertices are subsets of $[n]$ of size r , denoted $\binom{[n]}{r}$. In $J(n, r)$, two

¹The requirement that this operation be controlled on the first two registers, i.e., that it always leaves the first two registers unchanged, is not explicitly stated in [MNRS11]. However, this condition is needed to prevent, for example, the action $|x\rangle |\psi\rangle |0, 0\rangle \mapsto |x\rangle |D(x)\rangle |\phi\rangle$, where $\langle \psi | D(x) \rangle = 0$ and $|x\rangle |D(x)\rangle |\phi\rangle$ is a possible state of the algorithm. In this case, $(\text{LOCAL DIFFUSION})_{\text{ref}[0,0]} (\text{LOCAL DIFFUSION})^\dagger$ would not act as $W(P)$ on $|x\rangle |D(x)\rangle |\phi\rangle$.

vertices S, S' are adjacent if $|S \cap S'| = r - 1$. The data function is $D(S) = \{(i, \chi_i) : i \in S\}$. The diffusion step of this walk acts as

$$|S\rangle |D(S)\rangle |0\rangle \mapsto |S\rangle |D(S)\rangle \frac{1}{\sqrt{r(n-r)}} \sum_{i \in S, j \in [n] \setminus S} |(S \setminus i) \cup j\rangle |D((S \setminus i) \cup j)\rangle.$$

We can perform this diffusion in two queries by performing the transformation

$$|S\rangle |D(S)\rangle |0\rangle \mapsto |S\rangle |D(S)\rangle \frac{1}{\sqrt{r}} \sum_{i \in S} |(i, \chi_i)\rangle \frac{1}{\sqrt{n-r}} \sum_{j \in [n] \setminus S} |(j, \chi_j)\rangle.$$

We can reversibly map this to the desired state with no queries, and by using an appropriate encoding of D , we can make this time efficient as well.

To complete the description of this algorithm, we describe the marked set and checking procedure. We deviate slightly from the usual quantum walk algorithm of [Amb04] and instead describe a variation that is analogous to the learning graph for Element Distinctness [Bel12b]. We say a vertex S is marked if it contains an index i such that there exists $j \in [n] \setminus \{i\}$ with $\chi_i = \chi_j$ (whereas in [Amb04] both i and j must be in S). To check if S is marked, we simply search over $[n] \setminus S$ for such a j , in cost $O(\sqrt{n})$. This does not give asymptotically better performance than [Amb04], but it is more analogous to the 3-Distinctness algorithm we attempt to construct in the remainder of this section, and then succeed in constructing in Section 4.

Attempting a Quantum Walk for 3-Distinctness. We now attempt to construct an analogous algorithm for 3-Distinctness. Conceptually, the approach is simple, but successfully implementing the simple idea is nontrivial. The idea is to walk on a Johnson graph of sets of *collision pairs*, analogous to the set of queried indices in the Element Distinctness walk described above. The checking step is then similar to that of the above walk: simply search for a third element that forms a 3-collision with one of the 2-collisions in the set. For the update step, we need to replace one of the collision pairs in the set using a subroutine that finds a 2-collision. We now describe the difficulty of implementing this step efficiently, despite having an optimal Element Distinctness algorithm at our disposal. Section 3 presents a framework that allows us to successfully implement the idea in Section 4.

Let \mathcal{P} denote the set of collision pairs in the input, and $n_2 = |\mathcal{P}|$. We walk on $J(n_2, s_2)$, with each vertex S_2 corresponding to a set of s_2 collision pairs. The diffusion for this walk is the map

$$|S_2, D(S_2)\rangle |0\rangle \mapsto |S_2, D(S_2)\rangle \frac{1}{\sqrt{r(n_2-s_2)}} \sum_{\substack{(i,i') \in S_2 \\ (j,j') \in \mathcal{P} \setminus S_2}} |(S_2 \setminus (i, i')) \cup (j, j')\rangle |D((S_2 \setminus (i, i')) \cup (j, j'))\rangle.$$

To accomplish this, we need to generate $\frac{1}{\sqrt{s_2}} \sum_{(i,i') \in S_2} |(i, i', \chi_i)\rangle$ and $\frac{1}{\sqrt{n_2-s_2}} \sum_{(j,j') \in \mathcal{P} \setminus S_2} |(j, j', \chi_j)\rangle$. The first superposition is easy to generate, since we have S_2 , but the second is more difficult since we have to find new collisions.

The obvious approach is to use the quantum walk algorithm for Element Distinctness as a subroutine. However, this algorithm does not return the desired superposition over collisions; rather, it returns a superposition over sets that contain a collision. That is, we have the state $\frac{1}{\sqrt{n_2}} \sum_{(i,i') \in \mathcal{P}} |(i, i', \chi_i)\rangle |\psi(i, i')\rangle$ for some garbage $|\psi(i, i')\rangle$. The garbage may be only slightly entangled with (i, i') , but even this small amount of error in the state is prohibitive. Since we must call the update subroutine many times, we need the error to be very small. Unlike for nested checking, where bounded-error subroutines are sufficient, we cannot amplify the success probability of an update operator. We cannot directly use the state returned by the Element Distinctness

algorithm for several reasons. First, we cannot append garbage each time we update, as this would prevent proper interference in the walk. Second, when we use a nested walk for the update step, we would like to use the same trick as in nested checking: putting a copy of the starting state for the nested walk in the data structure so that we only need to perform the inner setup once. To do the same here, we would need to preserve the inner walk starting state; in other words, the update would need to output some state close to $\binom{n}{s_1}^{-1/2} \sum_{S_1 \in \binom{[n]}{s_1}} |S_1\rangle$. While we might try to recycle the garbage to produce this state, it is unclear how to extract the part we need for the update coherently, let alone without damaging the rest of the state.

This appears to be a problem for any approach that directly uses a quantum walk for the update, since all known quantum walks use some variant of a Johnson graph. Our modified framework circumvents this issue by allowing us to do the update with some garbage, which we then uncompute. This lets us use a quantum walk subroutine, with setup performed only at the beginning of the algorithm, to accomplish the update step. More generally, using our modified framework, we can tolerate updates that have garbage for any reason, whether the garbage is the result of the update being implemented by a quantum walk, or by some other quantum subroutine.

3 Quantum Walks with Nested Updates

3.1 Coin-Dependent Data

A quantum analog of a discrete-time random walk on a graph can be constructed as a unitary process on the directed edges. For an edge $\{x, y\}$, we may have a state $|x\rangle |y\rangle$, where $|x\rangle$ represents the current vertex and $|y\rangle$ represents the *coin* or next vertex. In the framework of [MNRS11], some data function on the vertices is employed to help implement the search algorithm. We modify the quantum walk framework to allow this data to depend on both the current vertex and the coin, so that it is a function of the directed edges, which seems natural in hindsight. We show that this point of view has algorithmic applications. In particular, this modification enables efficient nested updates.

In the rest of the paper, let P be a reversible, ergodic Markov chain on a connected, undirected graph $G = (X, E)$ with stationary distribution π and spectral gap $\delta > 0$.

Let $0 \notin X$. Let $D : (X \times \{0\}) \cup \vec{E} \rightarrow \mathcal{D}$ for some Hilbert space \mathcal{D} . A quantum analog of P with coin-dependent data structures can be implemented using three operations, as in [MNRS11], but the update now has three parts. The first corresponds to LOCAL DIFFUSION from the framework of [MNRS11], as described in Section 2.1. The others are needed because of the new coin-dependent data.

Update cost: Let U be the cost of implementing

- LOCAL DIFFUSION: $|x, 0\rangle |D(x, 0)\rangle \mapsto \sum_{y \in X} \sqrt{P(x, y)} |x, y\rangle |D(x, y)\rangle \forall x \in X$;
- The $(X, 0)$ -PHASE FLIP: $|x, 0\rangle |D(x, 0)\rangle \mapsto -|x, 0\rangle |D(x, 0)\rangle \forall x \in X$, and the identity on the orthogonal subspace; and
- The DATABASE SWAP: $|x, y\rangle |D(x, y)\rangle \mapsto |y, x\rangle |D(y, x)\rangle \forall (x, y) \in \vec{E}$.

By cost, we mean any desired measure of complexity such as queries, time, or space. We also naturally extend the setup and checking costs as follows, where $M \subseteq X$ is a set of marked vertices.

Setup cost: Let S be the cost of constructing

$$|\pi\rangle := \sum_{x \in X} \sqrt{\pi(x)} \sum_{y \in X} \sqrt{P(x, y)} |x, y\rangle |D(x, y)\rangle.$$

Checking cost: Let C be the cost of the reflection

$$|x, y\rangle |D(x, y)\rangle \mapsto \begin{cases} -|x, y\rangle |D(x, y)\rangle & \text{if } x \in M, \\ |x, y\rangle |D(x, y)\rangle & \text{otherwise,} \end{cases} \quad \forall (x, y) \in \vec{E}.$$

Observe that $|\pi\rangle^0 := \sum_{x \in X} \sqrt{\pi(x)} |x, 0\rangle |D(x, 0)\rangle$ can be mapped to $|\pi\rangle$ by the LOCAL DIFFUSION, which has cost $U < S$, so we can also consider S to be the cost of constructing $|\pi\rangle^0$.

Theorem 3.1. *Let P be a Markov chain on $G = (X, E)$ with spectral gap $\delta > 0$, and let D be a coin-dependent data structure for P . Let $M \subseteq X$ satisfy $\Pr_{x \sim \pi}(x \in M) \geq \varepsilon > 0$ whenever $M \neq \emptyset$. Then there is a quantum algorithm that finds an element of M , if $M \neq \emptyset$, with bounded error and with cost*

$$O\left(S + \frac{1}{\sqrt{\varepsilon}} \left(\frac{1}{\sqrt{\delta}} U + C\right)\right).$$

Proof. Our quantum walk algorithm is nearly identical to that of [MNRS11], so the proof of this theorem is also very similar. Just as in [MNRS11], we define a walk operator, $W(P)$, and analyze its spectral properties. Let $\mathcal{A} := \text{span}\{\sum_y \sqrt{P(x, y)} |x, y\rangle |D(x, y)\rangle : x \in X\}$ and define $W(P) := ((\text{DATABASE SWAP}) \cdot \text{ref}_{\mathcal{A}})^2$, where $\text{ref}_{\mathcal{A}}$ denotes the reflection about \mathcal{A} .

As in [MNRS11], we can define $\mathcal{H} := \text{span}\{|x, y\rangle : (x, y) \in (X \times \{0\}) \cup \vec{E}\}$ and $\mathcal{H}_D := \text{span}\{|x, y, D(x, y)\rangle : (x, y) \in (X \times \{0\}) \cup \vec{E}\}$. As in [MNRS11], there is a natural isomorphism $|x, y\rangle \mapsto |x, y\rangle_D = |x, y, D(x, y)\rangle$, and \mathcal{H}_D is invariant under both $W(P)$ and the checking operation. Thus, the spectral analysis may be done in \mathcal{H} , on states without data, *exactly* as in [MNRS11]. However, there are some slight differences in how we implement $W(P)$, which we now discuss.

The first difference is easy to see: in [MNRS11], the DATABASE SWAP can be accomplished trivially by a SWAP operation, mapping $|x\rangle |y\rangle |D(x)\rangle |D(y)\rangle$ to $|y\rangle |x\rangle |D(y)\rangle |D(x)\rangle$, whereas in our case, there may be a nontrivial cost associated with the mapping $|D(x, y)\rangle \mapsto |D(y, x)\rangle$, which we must include in the calculation of the update cost.

The second difference is more subtle. In [MNRS11], $\text{ref}_{\mathcal{A}}$ is implemented by applying $(\text{LOCAL DIFFUSION})^\dagger$, reflecting about $|0, D(0)\rangle$ (since the data only refers to a vertex) in the coin register, and then applying (LOCAL DIFFUSION) . It is simple to reflect about $|0, D(0)\rangle$, since $|D(0)\rangle = |0\rangle$ in the formalism of [MNRS11]. In [MNRS11], this reflection is sufficient, because the operation $(\text{LOCAL DIFFUSION})^\dagger$ fixes the vertex and its data, $|x\rangle |D(x)\rangle$, so in particular, it is still in the space $\text{span}\{|x\rangle |D(x)\rangle : x \in X\}$. The register containing the coin and its data, $|y\rangle |D(y)\rangle$, may be moved out of this space by $(\text{LOCAL DIFFUSION})^\dagger$, so we must reflect about $|0\rangle |D(0)\rangle$, but this is straightforward.

With coin-dependent data, a single register $|D(x, 0)\rangle$ holds the data for both the vertex and its coin, and the operation $(\text{LOCAL DIFFUSION})^\dagger$ may take the coin as well as the entire data register out of the space \mathcal{H}_D , so we need to reflect about $|0\rangle |D(x, 0)\rangle$, which is not necessarily defined to be $|0\rangle |0\rangle$. This explains why the cost of $(X, 0)$ -PHASE FLIP is also part of the update cost. In summary, we implement $W(P)$ by $((\text{DATABASE SWAP}) \cdot (\text{LOCAL DIFFUSION}) \cdot ((X, 0)\text{-PHASE FLIP}) \cdot (\text{LOCAL DIFFUSION})^\dagger)^2$. \square

3.2 Nested Updates

We show how to implement efficient nested updates using the coin-dependent data framework. Let $C : X \cup \{0\} \rightarrow \mathcal{C}$ be some coin-independent data structure (that will be a part of the final data structure) with $|C(0)\rangle = |0\rangle$, where we can reflect about $\text{span}\{|x\rangle |C(x)\rangle : x \in M\}$ in cost C_C . In the motivating example, if $x = S_2$ is a set of collision pairs, then $C(S_2)$ stores their query values.

Fix $x \in X$. Let P^x be a walk on a graph $G^x = (V^x, E^x)$ with stationary distribution π^x and marked set $M^x \subset V^x$. We use this walk to perform LOCAL DIFFUSION over $|x\rangle$. Let d^x be the data for this walk.

When there is ambiguity, we specify the data structure with a subscript. For instance, $|\pi\rangle_D = \sum_{x,y \in X} \sqrt{\pi(x)P(x,y)} |x,y\rangle |D(x,y)\rangle$ and $|\pi\rangle_C^0 = \sum_{x \in X} \sqrt{\pi(x)} |x,0\rangle |C(x),0\rangle$. Similarly, S_C is the cost to construct the state $|\pi\rangle_C$.

Definition 3.2. *The family $(P^x, M^x, d^x)_{x \in X}$ implements the LOCAL DIFFUSION and DATABASE SWAP of (P, C) with cost T if the following two maps can be implemented with cost T :*

LOCAL DIFFUSION WITH GARBAGE: *For some garbage states $(|\psi(x,y)\rangle)_{(x,y) \in \vec{E}}$, an operation controlled on the vertex x and $C(x)$, acting as*

$$|x,0\rangle |C(x),0\rangle |\pi^x(M^x)\rangle_{d^x} \mapsto \sum_{y \in X} \sqrt{P(x,y)} |x,y\rangle |C(x),C(y)\rangle |\psi(x,y)\rangle;$$

GARBAGE SWAP: *For any edge $(x,y) \in \vec{E}$,*

$$|x,y\rangle |C(x),C(y)\rangle |\psi(x,y)\rangle \mapsto |y,x\rangle |C(y),C(x)\rangle |\psi(y,x)\rangle.$$

The data structure of the implementation is $|D(x,0)\rangle = |C(x),0\rangle |\pi^x(M^x)\rangle_{d^x}$ for all $x \in X$ and $|D(x,y)\rangle = |C(x),C(y)\rangle |\psi(x,y)\rangle$ for any edge $(x,y) \in \vec{E}$.

Theorem 3.3. *Let P be a reversible, ergodic Markov chain on $G = (X, E)$ with spectral gap $\delta > 0$, and let C be a data structure for P . Let $M \subseteq X$ be such that $\Pr_{x \sim \pi}(x \in M) \geq \varepsilon$ for some $\varepsilon > 0$ whenever $M \neq \emptyset$. Let $(P^x, M^x, d^x)_{x \in X}$ be a family implementing the LOCAL DIFFUSION and DATABASE SWAP of (P, C) with cost T , and let $S', U', C', 1/\varepsilon', 1/\delta'$ be upper bounds on the costs and parameters associated with each of the (P^x, M^x, d^x) . Then there is a quantum algorithm that finds an element of M , if $M \neq \emptyset$, with bounded error and with cost*

$$\tilde{O}\left(S_C + S' + \frac{1}{\sqrt{\varepsilon}} \left(\frac{1}{\sqrt{\delta}} \left(\frac{1}{\sqrt{\varepsilon'}} \left(\frac{1}{\sqrt{\delta'}} U' + C' \right) + \mathsf{T} \right) + C_C \right)\right).$$

Proof. We achieve this upper bound using the quantization of P with the data structure of the implementation, D . We must compute the cost of the setup, update, and checking operations associated with this walk.

Checking: The checking cost $C = C_D$ is the cost to reflect about $\text{span}\{|x\rangle |y\rangle |D(x,y)\rangle : x \in M\} = \text{span}\{|x\rangle |y\rangle |C(x),C(y)\rangle |\psi(x,y)\rangle : x \in M\}$. We can implement this in \mathcal{H}_D by reflecting about $\text{span}\{|x\rangle |C(x)\rangle : x \in M\}$, which costs C_C .

Setup: Recall that $|C(0)\rangle = |0\rangle$. The setup cost $S = S_D$ is the cost of constructing the state

$$\sum_{x \in X} \sqrt{\pi(x)} |x\rangle |0\rangle |D(x,0)\rangle = \sum_{x \in X} \sqrt{\pi(x)} |x\rangle |0\rangle |C(x),0\rangle |\pi^x(M^x)\rangle.$$

We do this as follows. We first construct $\sum_{x \in X} \sqrt{\pi(x)} |x, 0\rangle |C(x), 0\rangle$ in cost S_C . Next, we apply the mapping $|x\rangle \mapsto |x\rangle |\pi^x\rangle$ in cost S' . Finally, we use the quantization of P^x to perform the mapping $|x\rangle |\pi^x\rangle \mapsto |x\rangle |\pi^x(M^x)\rangle$ in cost $\frac{1}{\sqrt{\varepsilon'}}(\frac{1}{\sqrt{\delta'}}U' + C')$. The full setup cost is then $S = S_C + S' + \frac{1}{\sqrt{\varepsilon'}}(\frac{1}{\sqrt{\delta'}}U' + C')$.

Update: The update cost has three contributions. The first is the LOCAL DIFFUSION operation, which, by the definition of D , is exactly the LOCAL DIFFUSION WITH GARBAGE operation. Similarly, the DATABASE SWAP is exactly the GARBAGE SWAP, so these two operations have total cost T . The $(X, 0)$ -PHASE FLIP is simply a reflection about states of the form $|x\rangle |D(x, 0)\rangle = |x\rangle |C(x)\rangle |\pi^x(M^x)\rangle$. Given any $x \in X$, we can reflect about $|\pi^x(M^x)\rangle$ using the quantization of P^x in cost $\frac{1}{\sqrt{\varepsilon'}}(\frac{1}{\sqrt{\delta'}}U' + C')$ by running the algorithm of Theorem 3.1. In particular, we can run the walk backward to prepare the state $|\pi^x\rangle$, perform phase estimation on the walk operator to implement the reflection about this state, and then run the walk forward to recover $|\pi^x(M^x)\rangle$. However, this transformation is implemented approximately. To keep the overall error small, we need an accuracy of $O(1/\sqrt{\varepsilon\delta\varepsilon'\delta'})$, which leads to an overhead logarithmic in the required accuracy. The reflection about $|\pi^x(M^x)\rangle$, controlled on $|x\rangle$, is sufficient because LOCAL DIFFUSION WITH GARBAGE is controlled on $|x\rangle |C(x)\rangle$, and so it leaves these registers unchanged. Since we apply the $(X, 0)$ -PHASE FLIP just after applying (LOCAL DIFFUSION) † (see proof of Theorem 3.1) to a state in \mathcal{H}_D , we can guarantee that these registers contain $|x\rangle |C(x)\rangle$ for some $x \in X$. The total update cost (up to log factors) is $U = T + \frac{1}{\sqrt{\varepsilon'}}(\frac{1}{\sqrt{\delta'}}U' + C')$.

Finally, the full cost of the quantization of P (up to log factors) is

$$\begin{aligned} S_C + S' + \frac{1}{\sqrt{\varepsilon'}} \left(\frac{1}{\sqrt{\delta'}} U' + C' \right) + \frac{1}{\sqrt{\varepsilon}} \left(\frac{1}{\sqrt{\delta}} \left(\frac{1}{\sqrt{\varepsilon'}} \left(\frac{1}{\sqrt{\delta'}} U' + C' \right) + T \right) + C_C \right) \\ = \tilde{O} \left(S_C + S' + \frac{1}{\sqrt{\varepsilon}} \left(\frac{1}{\sqrt{\delta}} \left(\frac{1}{\sqrt{\varepsilon'}} \left(\frac{1}{\sqrt{\delta'}} U' + C' \right) + T \right) + C_C \right) \right). \end{aligned} \quad \square$$

If $T = 0$ (as may be the case, e.g., when the notion of cost is query complexity), then the expression is exactly what we would have liked for nested updates.

4 Application: Quantum Query Complexity of 3-Distinctness

In this section we prove the following theorem.

Theorem 4.1. *The quantum query complexity of 3-Distinctness is $\tilde{O}(n^{5/7})$.*

We begin by giving a high-level description of the quantum walk algorithm before describing the implementation of each required procedure and their costs. First we define some notation.

For any set $S_1 \subseteq A_1 \cup A_2$, let $\mathcal{P}(S_1) := \{(i, j) \in A_1 \times A_2 : i, j \in S_1, i \neq j, \chi_i = \chi_j\}$ be the set of 2-collisions in S_1 and for any set $S_2 \subset A_1 \times A_2$, let $\mathcal{I}(S_2) := \bigcup_{(i, j) \in S_2} \{i, j\}$ be the set of indices that are part of pairs in S_2 . In general, we only consider 2-collisions in $A_1 \times A_2$; other 2-collisions in χ are ignored. For any pair of sets A, B , let $\mathcal{P}(A, B) := \{(i, j) \in A \times B : i \neq j, \chi_i = \chi_j\}$ be the set of 2-collisions between A and B . For convenience, we define $\mathcal{P} := \mathcal{P}(A_1, A_2)$. Let $n_2 := |\mathcal{P}|$ be the size of this set. For any set $S_2 \subseteq \mathcal{P}$, we denote the set of queried values by $Q(S_2) := \{(i, j, \chi_i) : (i, j) \in S_2\}$. Similarly, for any set $S_1 \subset [n]$, we denote the set of queried values by $Q(S_1) := \{(i, \chi_i) : i \in S_1\}$.

4.1 High-Level Description of the Walk

The Walk Our overall strategy is to find a 2-collision $(i, j) \in A_1 \times A_2$ such that $\exists k \in A_3$ with $\{i, j, k\}$ a 3-collision. Let $s_1, s_2 < n$ be parameters to be optimized. We walk on the vertices $X = \binom{\mathcal{P}}{s_2}$, with each vertex corresponding to a set of s_2 2-collisions from $A_1 \times A_2$. A vertex is considered marked if it contains (i, j) such that $\exists k \in A_3$ with $\{i, j, k\}$ a 3-collision. Thus, if $M \neq \emptyset$, the proportion of marked vertices is $\varepsilon = \Omega(\frac{s_2}{n_2})$.

To perform an update, we use an Element Distinctness subroutine that walks on s_1 -sized subsets of $A_1 \cup A_2$. However, since n_2 is large by assumption, the expected number of collisions in a set of size s_1 is large if $s_1 \gg \sqrt{n}$, which we suppose holds. It would be a waste to take only one and leave the rest, so we replace multiple elements of S_2 in each step. This motivates using a generalized Johnson graph $J(n_2, s_2, m)$ for the main walk, where we set $m := \frac{s_1^2 n_2}{n^2} = O(\frac{s_1^2}{n})$, the expected number of 2-collisions in a set of size s_1 . In $J(n_2, s_2, m)$, two vertices S_2 and S'_2 are adjacent if $|S_2 \cap S'_2| = s_2 - m$, so we can move from S_2 to S'_2 by replacing m elements of S_2 by m distinct elements. Let $\Gamma(S_2)$ denote the set of vertices adjacent to S_2 . The spectral gap of $J(n_2, s_2, m)$ is $\delta = \Omega(\frac{m}{s_2})$.

The Update To perform an update step on the vertex S_2 , we use the Element Distinctness algorithm of [Amb04] as a subroutine, with some difference in how we define the marked set. Specifically, we use the subroutine to look for m 2-collisions, with $m \gg 1$. Furthermore, we only want to find 2-collisions that are not already in S_2 , so P^{S_2} is a walk on $J(2n/3 - 2s_2, s_1)$, with vertices corresponding to sets of s_1 indices from $(A_1 \cup A_2) \setminus \mathcal{I}(S_2)$, and we consider a vertex marked if it contains at least m pairs of indices that are 2-collisions (i.e., $M^{S_2} = \{S_1 \in \binom{(A_1 \cup A_2) \setminus \mathcal{I}(S_2)}{s_1} : |\mathcal{P}(S_1)| \geq m\}$).

The Data We store the value χ_i with each $(i, j) \in S_2$ and $i \in S_1$, i.e., $|C(S_2)\rangle = |Q(S_2)\rangle$ and $|d^{S_2}(S_1, S'_1)\rangle = |Q(S_1), Q(S'_1)\rangle$. Although technically this is part of the data, it is classical and coin-independent, so it is straightforward. Furthermore, since S_1 is encoded in $Q(S_1)$ and S_2 in $Q(S_2)$, we simply write $|Q(S_1)\rangle$ instead of $|S_1, Q(S_1)\rangle$ and $|Q(S_2)\rangle$ instead of $|S_2, Q(S_2)\rangle$.

The rest of the data is what is actually interesting. We use the state $|\pi^{S_2}(M^{S_2})\rangle_{d^{S_2}}^0$ in the following instead of $|\pi^{S_2}(M^{S_2})\rangle_{d^{S_2}}$ since it is easy to map between these two states. For every $S_2 \in X$, let

$$|D(S_2, 0)\rangle := |Q(S_2), 0\rangle |\pi^{S_2}(M^{S_2})\rangle_{d^{S_2}}^0 = |Q(S_2)\rangle \frac{1}{\sqrt{|M^{S_2}|}} \sum_{S_1 \in M^{S_2}} |Q(S_1)\rangle,$$

and for every edge (S_2, S'_2) , let $|D(S_2, S'_2)\rangle := |Q(S_2), Q(S'_2)\rangle |\psi(S_2, S'_2)\rangle$ where

$$|\psi(S_2, S'_2)\rangle := \sum_{\tilde{S}_1 \in \binom{(A_1 \cup A_2) \setminus \mathcal{I}(S_2 \cup S'_2)}{s_1 - 2m}} \sqrt{\frac{\binom{n_2 - s_2}{m}}{(|\mathcal{P}(\tilde{S}_1)| + m) |M^{S_2}|}} |Q(\tilde{S}_1)\rangle. \quad (1)$$

We define $|\psi\rangle$ in this way precisely because it is what naturally occurs when we attempt to perform the diffusion.

4.2 Implementation and Cost Analysis

We now explain how to implement the walk described at the beginning of this section and analyze the costs of the associated operations.

We have assumed that we have some partition A_1, A_2, A_3 of $[n]$, although we actually want to run our algorithm on a random partition. The starting state is a uniform superposition over s_2 collision pairs across the bipartition $A_1 \times A_2$. Unfortunately, given A_1, A_2 , we are unable to construct a valid starting state. However, we can generate a state-partition pair $(|\pi(A_1, A_2)\rangle, A_1, A_2)$ such that the distribution of $A_1, A_2, A_3 := [n] \setminus (A_1 \cup A_2)$ is sufficiently random, and $|\pi(A_1, A_2)\rangle$ is a starting state for the partition A_1, A_2, A_3 .

Theorem 4.2 (Outer walk setup cost S_C). *The starting state of the outer walk, $\binom{n_2}{s_2}^{-1/2} \sum_{S_2 \in \binom{\mathcal{P}(A_1, A_2)}{s_2}} |Q(S_2)\rangle$, can be constructed for random variables A_1, A_2, A_3 with $|A_1| = |A_2| = |A_3| = n/3$, such that if χ has a unique 3-collision $\{i, j, k\}$, then $\Pr((i, j, k) \in A_1 \times A_2 \times A_3) = \Omega(1)$, in $\tilde{O}(s_1 + s_2 \sqrt{n/s_1})$ queries.*

Proof. To begin, we choose a random tripartition $\tilde{A}_1, \tilde{A}_2, \tilde{A}_3$ of $[n]$ such that $|\tilde{A}_1| = \frac{n}{3} + s_1 - s_2$, $|\tilde{A}_2| = \frac{n}{3}$, and $|\tilde{A}_3| = \frac{n}{3} - s_1 + s_2$. Our final sets A_1, A_2, A_3 are closely related to these sets, but satisfy $|A_1| = |A_2| = |A_3| = \frac{n}{3}$. Let \tilde{n}_2 be the number of 2-collisions across $\tilde{A}_1 \times \tilde{A}_2$. We first create a uniform superposition over all subsets of \tilde{A}_1 of size s_1 along with their query values, $\binom{n/3+s_1-s_2}{s_1}^{-1/2} \sum_{I \in \binom{\tilde{A}_1}{s_1}} |Q(I)\rangle$, using $O(s_1)$ queries.

For a set $I \in \binom{\tilde{A}_1}{s_1}$, let $H(I) \subset \tilde{A}_2$ denote the set $\{j \in \tilde{A}_2 : \exists i \in I, \chi_i = \chi_j\}$ of indices in \tilde{A}_2 colliding with I . Next we repeatedly Grover search for indices in $H(I)$. For a uniform I , the size of $H(I)$ is roughly $\frac{\tilde{n}_2 s_1}{n} = \Omega(s_1)$ in expectation; more specifically, for most choices \tilde{A}_1 and \tilde{A}_2 , we have $\Pr_I(|H(I)| \in \Omega(s_1)) \geq 1 - o(1)$. We can therefore consider only the part of the state $\binom{n/3+s_1-s_2}{s_1}^{-1/2} \sum_{I \in \binom{\tilde{A}_1}{s_1}: |H(I)| \geq \epsilon s_1} |Q(I)\rangle$, for a suitable constant ϵ . Thus, we can use Grover search to find and query s_2 elements of $H(I)$ in $\tilde{O}(s_2 \sqrt{n/s_1})$ queries, obtaining a state close to

$$\binom{n/3 + s_1 - s_2}{s_1}^{-1/2} \sum_{I \in \binom{\tilde{A}_1}{s_1}: |H(I)| \geq \epsilon s_1} |Q(I)\rangle \binom{|H(I)|}{s_2}^{-1/2} \sum_{J \in \binom{H(I)}{s_2}} |Q(J)\rangle.$$

For a given J , we can partition the set I into two disjoint sets: I_1 , which contains all elements in I that do not collide with any element in J ; and I_2 , which contains elements that do collide with an element in J . We can then combine I_2 with J to get a set of s_2 collision pairs. The full reversible mapping, which costs 0 queries, is $|Q(I), Q(J)\rangle \mapsto |Q(I_1)\rangle |\{(i, j, \chi_i) : i \in I_2, j \in J\}\rangle$. Applying this transformation gives (a state close to)

$$\binom{n/3 + s_1 - s_2}{s_1}^{-1/2} \sum_{I_1 \in \binom{\tilde{A}_1}{s_1 - s_2}: |H(I_1)| \geq \epsilon s_1 - s_2} |Q(I_1)\rangle \binom{|H(I_1)| + s_2}{s_2}^{-1/2} \sum_{S_2 \in \mathcal{P}(\tilde{A}_1 \setminus I_1, \tilde{A}_2)} |Q(S_2)\rangle.$$

Note that this state is not uniform in I_1 , but is uniform in S_2 when we restrict to a particular I_1 . Thus we measure the first register to get some I_1 with non-uniform probability that depends only

on $|H(I_1)|$. The remaining state is the uniform superposition

$$\left(|\mathcal{P}(\tilde{A}_1 \setminus I_1, \tilde{A}_2)| \right)^{-1/2} \sum_{S_2 \in \binom{\mathcal{P}(\tilde{A}_1 \setminus I_1, \tilde{A}_2)}{s_2}} |Q(S_2)\rangle.$$

Now let $A_1 = \tilde{A}_1 \setminus I_1$, $A_2 = \tilde{A}_2$ and $A_3 = \tilde{A}_3 \cup I_1$. Then we have $\mathcal{P} = \mathcal{P}(\tilde{A}_1 \setminus I_1, \tilde{A}_2) = \mathcal{P}(A_1, A_2)$, so we have constructed the correct state for the tripartition A_1, A_2, A_3 . Clearly, if $\{i, j, k\}$ is the unique 3-collision, then $i \in \tilde{A}_1$, $j \in \tilde{A}_2$ and $k \in \tilde{A}_3$ with constant probability. It remains to consider whether $i \in I_1$. Although the distribution of I_1 is non-uniform, the distribution restricted to those I_1 with $H(I_1) = h$ is uniform for any fixed h , and it is easy to see that $\Pr(i \in I_1 | H(I_1) = h)$ is $o(1)$ for any h .

For more details, refer to the proof of Theorem 5.4, which also proves an analogous statement for time complexity. \square

Hereafter, we assume the above choice of partition A_1, A_2, A_3 , and that if there is a unique 3-collision $\{i, j, k\}$, then $i \in A_1$, $j \in A_2$ and $k \in A_3$.

Theorem 4.3 (Costs of the update walk S' , $\frac{1}{\sqrt{\varepsilon'}}(\frac{1}{\sqrt{\delta'}}U' + C')$). *The update walk has query complexities $S' = O(s_1)$ and $\frac{1}{\sqrt{\varepsilon'}}(\frac{1}{\sqrt{\delta'}}U' + C') = \tilde{O}(\sqrt{nm/s_1})$.*

Proof. Fix an arbitrary vertex $S_2 \in \binom{\mathcal{P}(A_1, A_2)}{s_2}$. We now analyze the update walk P^{S_2} . The walk is still on $J(2n/3 - 2s_2, s_1)$, so $\delta' = \Omega(\frac{1}{s_1})$, but in contrast to [Amb04], a vertex is considered marked if it has at least m collision pairs, and we have a lower bound of $n_2 = \Omega(n)$ on the number of disjoint collision pairs. Since we defined $m = s_1^2 n_2 / n^2$ as roughly the expected number of collision pairs in a set of size s_1 , we have $\varepsilon' = \Omega(1)$. We still need to do the walk, both to amplify the success probability to inverse polynomial (which we could also have done by increasing s_1 by log factors) and more importantly, to implement the phase flip $|x, 0\rangle |D(x, 0)\rangle \mapsto -|x, 0\rangle |D(x, 0)\rangle$.

Setup: We need $S' = O(s_1)$ queries to set up $\left(\binom{2n/3 - 2s_2}{s_1} \right)^{-1/2} \sum_{S_1 \in \binom{(A_1 \cup A_2) \setminus \mathcal{I}(S_2)}{s_1}} |Q(S_1)\rangle$.

Update: The update on $J(2n/3 - 2s_2, s_1)$ costs $O(1)$ queries.

Checking: The query complexity of checking is 0, since we merely observe whether there are m colliding pairs in S_1 .

We can thus compute $\frac{1}{\sqrt{\varepsilon'}}(\frac{1}{\sqrt{\delta'}}U' + C') = \tilde{O}(\sqrt{\frac{mn^2}{s_1^2 n_2}} \sqrt{s_1}) = \tilde{O}(\sqrt{\frac{nm}{s_1}})$. \square

The following lemma tells us that we do not need to reverse the garbage part of the data.

Lemma 4.4. *For all edges (S_2, S'_2) , $|\psi(S_2, S'_2)\rangle = |\psi(S'_2, S_2)\rangle$.*

Proof. Recall the definition of $|\psi(S_2, S'_2)\rangle$ from (1):

$$|\psi(S_2, S'_2)\rangle = \sum_{\tilde{S}_1 \in \binom{(A_1 \cup A_2) \setminus \mathcal{I}(S_2 \cup S'_2)}{s_1 - 2m}} \sqrt{\frac{\binom{n_2 - s_2}{m}}{\binom{|\mathcal{P}(\tilde{S}_1)| + m}{m} |M_{S_2}|}} |\psi(\tilde{S}_1)\rangle.$$

To see that this is symmetric in S_2 and S'_2 , we need only show that $|M^{S_2}| = |M^{S'_2}|$. We have

$$|M^{S_2}| = \left| \left\{ S_1 \in \binom{(A_1 \cup A_2) \setminus \mathcal{I}(S_2)}{s_1} : |\mathcal{P}(S_1)| \geq m \right\} \right| = \binom{|\mathcal{P} \setminus S_2|}{m} \binom{|(A_1 \cup A_2) \setminus \mathcal{I}(S_2)| - 2m}{s_1 - 2m}.$$

This holds because all collisions in $A_1 \cup A_2$ are disjoint, and so choosing m pairs from $\mathcal{P} \setminus S_2$ gives $2m$ distinct indices. We can easily see that $|\mathcal{P} \setminus S_2| = n_2 - s_2$, which is independent of S_2 . Less trivially, since all collisions in S_2 are disjoint, we have $|\mathcal{I}(S_2)| = 2s_2$ for all S_2 , and so $|(A_1 \cup A_2) \setminus \mathcal{I}(S_2)| = |(A_1 \cup A_2)| - 2s_2$, again, independent of S_2 . Thus we have $|M^{S_2}| = |M^{S'_2}|$, completing the proof. \square

From this lemma it readily follows that the GARBAGE SWAP requires no queries.

Theorem 4.5 (GARBAGE SWAP cost). *No queries are needed to perform the GARBAGE SWAP, which for any (S_2, S'_2) performs the map*

$$|Q(S_2), Q(S'_2)\rangle |\psi(S_2, S'_2)\rangle \mapsto |Q(S'_2), Q(S_2)\rangle |\psi(S'_2, S_2)\rangle.$$

The LOCAL DIFFUSION WITH GARBAGE also requires no queries, but is nontrivial to implement.

Theorem 4.6 (LOCAL DIFFUSION WITH GARBAGE cost). *No queries are needed to perform the LOCAL DIFFUSION WITH GARBAGE, which, for any S_2 , performs the map*

$$|Q(S_2)\rangle |\pi^{S_2}(M^{S_2})\rangle^0 \mapsto \frac{1}{\sqrt{|\Gamma(S_2)|}} \sum_{S'_2 \in \Gamma(S_2)} |Q(S_2), Q(S'_2)\rangle |\psi(S_2, S'_2)\rangle,$$

where $|\psi(S_2, S'_2)\rangle$ is defined in (1).

Proof. We employ the following procedure to perform the LOCAL DIFFUSION WITH GARBAGE.

1. Perform $|Q(S_2), Q(S_1)\rangle \mapsto |Q(S_2), Q(S_1)\rangle \binom{s_2}{m}^{-1/2} \sum_{I \in \binom{S_2}{m}} |Q(I)\rangle$.
2. Perform $|Q(S_2), Q(S_1)\rangle \mapsto |Q(S_2), Q(S_1)\rangle \binom{|\mathcal{P}(S_1)|}{m}^{-1/2} \sum_{J \in \binom{\mathcal{P}(S_1)}{m}} |Q(J)\rangle$.
3. Perform $|Q(S_1)\rangle |Q(J)\rangle \mapsto |Q(\tilde{S}_1)\rangle |Q(J)\rangle$, where $\tilde{S}_1 = S_1 \setminus \mathcal{I}(J)$.

Here I represents collision pairs to be removed from S_2 and J represents collision pairs to be added. It is clear that each of these operations has query complexity 0.

Now we show the correctness of this procedure. Recall that $S_1 \in M^{S_2}$ if and only if S_1 contains m collisions, i.e., $|\mathcal{P}(S_1)| \geq m$, so $|Q(S_2)\rangle |\pi^{S_2}(M^{S_2})\rangle^0$ is

$$|Q(S_2)\rangle \frac{1}{\sqrt{|M^{S_2}|}} \sum_{S_1 \in M^{S_2}} |Q(S_1)\rangle = |Q(S_2)\rangle \frac{1}{\sqrt{|M^{S_2}|}} \sum_{S_1 \in \binom{(A_1 \cup A_2) \setminus \mathcal{I}(S_2)}{s_1} : |\mathcal{P}(S_1)| \geq m} |Q(S_1)\rangle.$$

After performing the above procedure, we get the state

$$|Q(S_2)\rangle \frac{1}{\sqrt{|M^{S_2}|}} \sum_{S_1 \in \binom{(A_1 \cup A_2) \setminus \mathcal{I}(S_2)}{s_1} : |\mathcal{P}(S_1)| \geq m} |Q(\tilde{S}_1)\rangle \frac{1}{\sqrt{\binom{s_2}{m}}} \sum_{I \in \binom{S_2}{m}} |Q(I)\rangle \frac{1}{\sqrt{\binom{|\mathcal{P}(S_1)|}{m}}} \sum_{J \in \binom{\mathcal{P}(S_1)}{m}} |Q(J)\rangle.$$

Note that $\mathcal{P}(S_1) = \mathcal{P}(\tilde{S}_1) \cup J$. To see this, we must appeal to the fact that all collisions in $A_1 \times A_2$ are disjoint, by assumption, so for each collision pair $(i, j) \in J$, removing i and j from S_1 removes the collision pair (i, j) and no other collision pair from $\mathcal{P}(S_1)$. Next, we can see that $|\mathcal{P}(\tilde{S}_1) \cup J| = |\mathcal{P}(\tilde{S}_1)| + m$, since $J \cap \mathcal{P}(S_1) = \emptyset$ and $|J| = m$. Thus, $|\mathcal{P}(S_1)| = |\mathcal{P}(\tilde{S}_1)| + m$, and we can rewrite the state as

$$|Q(S_2)\rangle \binom{s_2}{m}^{-1/2} \binom{n_2 - s_2}{m}^{-1/2} \sum_{\substack{I \in \binom{S_2}{m} \\ J \in \binom{\mathcal{P}(S_2)}{m}}} |Q(I)\rangle |Q(J)\rangle \sum_{\tilde{S}_1 \in \binom{(A_1 \cup A_2) \setminus \mathcal{I}(S_2 \cup J)}{s_1 - 2m}} \alpha_{\tilde{S}_1}(S_2, (S_2 \cup J) \setminus I) |Q(\tilde{S}_1)\rangle,$$

where

$$\alpha_{\tilde{S}_1}(S_2, (S_2 \cup J) \setminus I) = \sqrt{\frac{\binom{n_2 - s_2}{m}}{(|\mathcal{P}(\tilde{S}_1)| + m) |M^{S_2}|}}.$$

We now simply note that the neighbours of any $S_2 \in X$ are exactly $(S_2 \cup J) \setminus I$ for $I \in \binom{S_2}{m}$ and $J \in \binom{\mathcal{P}(S_2)}{m}$. Furthermore, for such a neighbour $S'_2 = (S_2 \cup J) \setminus I$, $Q(S_2), Q(I), Q(J)$ encodes $Q(S_2), Q(S'_2)$. Finally, for such an S'_2 , we have $S_2 \cup J = S_2 \cup S'_2$. Thus, we are left with the desired state. \square

Corollary 4.7 (LOCAL DIFFUSION and DATABASE SWAP cost T). *The family $(P^{S_2}, M^{S_2}, d^{S_2})_{S_2 \in X}$ implements the LOCAL DIFFUSION and DATABASE SWAP of (P, Q) with no queries.*

Proof. This is immediate from Theorems 4.5 and 4.6. \square

The checking cost is immediate, since we can use Grover search to look for an element of A_3 that collides with any of the stored 2-collisions.

Theorem 4.8 (Checking cost C). *We can implement the checking reflection with $\mathsf{C} = \tilde{O}(\sqrt{n})$ queries.*

We now have all necessary ingredients to prove the main theorem.

Proof of Theorem 4.1. We apply Theorem 3.3 to compute the cost of our nested-update quantum walk algorithm, giving (up to log factors)

$$\begin{aligned} \mathsf{S}_C + \mathsf{S}' + \frac{1}{\sqrt{\varepsilon}} \left(\frac{1}{\sqrt{\delta}} \left(\frac{1}{\sqrt{\varepsilon'}} \left(\frac{1}{\sqrt{\delta'}} U' + \mathsf{C}' \right) + \mathsf{T} \right) + \mathsf{C} \right) \\ = s_1 + s_2 \sqrt{\frac{n}{s_1}} + s_1 + \sqrt{\frac{n_2}{s_2}} \left(\sqrt{\frac{s_2}{m}} \left(\frac{\sqrt{nm}}{\sqrt{s_1}} + 0 \right) + \sqrt{n} \right) \\ = s_1 + s_2 \sqrt{\frac{n}{s_1}} + \sqrt{\frac{n_2 n}{s_1}} + \sqrt{\frac{n_2 n}{s_2}} = \tilde{O} \left(s_1 + s_2 \sqrt{\frac{n}{s_1}} + \frac{n}{\sqrt{s_1}} + \frac{n}{\sqrt{s_2}} \right) \end{aligned}$$

using the cost calculations from Theorems 4.2, 4.3, and 4.8 and Corollary 4.7. Setting $s_1 = n^{5/7}$ and $s_2 = n^{4/7}$ gives query complexity $\tilde{O}(n^{5/7})$. \square

5 Time Complexity of 3-Distinctness

In this section we prove the following theorem.

Theorem 5.1. *The time complexity of 3-Distinctness is $\tilde{O}(n^{5/7})$.*

This follows fairly straightforwardly from the quantum walk described in Section 4. The only remaining task is to describe how we can encode the sets of queried indices and pairs of indices so that all necessary operations, such as inserting an element in a set or removing an element from a set, can be done in poly-logarithmic time. We use the same data structure that was used to obtain a tight upper bound on the time complexity of Element Distinctness [Amb04]. After describing the necessary properties of this data structure and how we apply it to our walk, we explain how each of the operations described in Section 4 can be done time-efficiently using this encoding.

The following lemma describes properties of the data structure that we use to encode edges of our walk and their data. We refer to this data structure as a *skip-list*.

Lemma 5.2 ([Amb04]). *There exists a data structure for storing a set of items of the form (z, χ) (the χ values need not be unique) that allows the following operations to be performed in worst case time complexity $O(\log^4(n + q))$: insert an item; delete an item; look up an item by its χ value; or create a superposition of the elements stored. The data structure storing a set S is a unique encoding of S .*

Encoding an Edge and its Data We now describe how to encode an edge and its data. These states have the form either $|S_2, S'_2, D(S_2, S'_2)\rangle$, where $|D(S_2, S'_2)\rangle$ is a superposition over basis states $|Q(S_2), Q(S'_2), Q(\tilde{S}_1)\rangle$ for $(S_2, S'_2) \in E$ and $\tilde{S}_1 \subset (A_1 \cup A_2) \setminus \mathcal{I}(S_2 \cup S'_2)$ (and recall that $Q(S_2), Q(S'_2)$ automatically encodes S_2, S'_2); or $|S_2, 0, D(S_2, 0)\rangle$, where $|D(S_2, 0)\rangle$ is a superposition over basis states $|Q(S_2), Q(S_1), Q(S'_1)\rangle$ for $S_2 \in X$, and $(S_1, S'_1) \in E^{S_2} \cup (V^{S_2} \times \{0\})$. Strictly speaking, we previously defined $|D(S_2, 0)\rangle$ using $|\pi^{S_2}(M^{S_2})\rangle_{d^{S_2}}^0$ instead of $|\pi^{S_2}(M^{S_2})\rangle_{d^{S_2}}$, that is, as a superposition of $|Q(S_2), Q(S_1), 0\rangle$ for $S_2 \in X$, and $S_1 \in V^{S_2}$. However, we must also consider how to encode states of the nested update walk, which do not generally have 0 in the coin register.

We begin by encoding the triple of sets $(Q(S_2), Q(S_1), Q(S'_1))$. We store each of $Q(S_2)$ and $Q(S_1)$ in a skip-table. To store $Q(S'_1)$ for $S'_1 \neq 0$, we simply store both $Q(S_1) \setminus Q(S'_1)$ and $Q(S'_1) \setminus Q(S_1)$, each of which is a single queried index (i, χ_i) . This already encodes the three sets, but we add additional structure to speed up certain tasks. We store $Q(\mathcal{P}(S_1))$, the set of 2-collisions in S_1 , in another skip-table. We also keep a counter of the size of this set so that we can easily check whether S_1 is marked.

Now we describe how we encode the triple of sets $(Q(S_2), Q(S'_2), Q(\tilde{S}_1))$. We store each of $Q(S_2)$ and $Q(\tilde{S}_1)$ in a skip-table. To store $Q(S'_2)$, we store each of $Q(S_2) \setminus Q(S'_2)$ and $Q(S'_2) \setminus Q(S_2)$ in a skip-table. We also store with $Q(\tilde{S}_1)$ an additional skip-table containing $Q(\mathcal{P}(\tilde{S}_1))$, with a counter encoding its size. We store this simply because this is what is left over from the encoding of $Q(S_1)$ after we perform LOCAL DIFFUSION.

It is now clear from Lemma 5.2 that we can perform the following operations in poly-logarithmic time: insert an element to S_1 , insert an element to S_2 , delete an element from S_1 , delete an element from S_2 , look up an element in S_1 , and look up an element in S_2 . In addition, we can perform each of these operations in superposition.

Cost Analysis We now analyze the cost of all operations implemented in Section 4.

Since we are now concerned with time complexity, we need some efficient way to store and compute the partition A_1, A_2, A_3 . We use the following notion to efficiently represent a random subset of $[n]$.

Definition 5.3. A family F of functions $f : [n] \rightarrow [\ell]$ is said to be k -wise independent if for any distinct $i_1, \dots, i_k \in [n]$, the distribution $(f(i_1), \dots, f(i_k))$ is identical to the uniform distribution over $[\ell]^k$.

When $n = \ell$ is a prime power, a simple example of a k -wise independent functions is the family of all polynomials of degree $k - 1$ over the finite field $\text{GF}(n)$ [WC81]. Each such polynomial can be represented using $O(k \log n)$ bits and evaluated using $O(k)$ additions and multiplications over $\text{GF}(n)$. More efficient constructions exist, but the polynomial construction suffices here. It can be extended to any integer n by allowing a small statistical distance from the distribution $(f(i_1), \dots, f(i_k))$ to the uniform distribution over $[n]^k$.

For simplicity, we now assume that we have at our disposal a (perfect) 3-wise independent family F of functions $f : [n] \rightarrow [n]$.

Theorem 5.4 (Outer walk setup cost S_C). *We can construct, in time $\tilde{O}(s_1 + s_2 \sqrt{n/s_1})$, a state $\binom{n_2}{s_2}^{-1/2} \sum_{S_2 \in \binom{\mathcal{P}(A_1, A_2)}{s_2}} |Q(S_2)\rangle$ for A_1, A_2 random variables such that*

1. $|A_1| = |A_2| = \frac{n}{3}$;
2. $A_1, A_2, A_3 := [n] \setminus (A_1 \cup A_2)$ is a tripartition of $[n]$;
3. if χ has a unique 3-collision $\{i, j, k\}$, $\Pr(i \in A_1, j \in A_2, k \in A_3) = \Omega(1)$;
4. the space complexity of storing the partition (A_1, A_2, A_3) is $\tilde{O}(s_1)$; and
5. the time complexity of determining to which of A_1, A_2 or A_3 an index $i \in [n]$ belongs is $\tilde{O}(1)$.

Proof. This proof is similar to that of Theorem 4.2, but we include significantly more detail. Let $f \in F$ be a 3-wise independent function, and define \tilde{A}_1, \tilde{A}_2 and \tilde{A}_3 by $f(i) \leq \frac{n}{3} + s_1 - s_2 \Leftrightarrow i \in \tilde{A}_1$, $\frac{n}{3} + s_1 - s_2 < f(i) \leq \frac{2n}{3} + s_1 - s_2 \Leftrightarrow i \in \tilde{A}_2$, and $f(i) > \frac{2n}{3} + s_1 - s_2 \Leftrightarrow i \in \tilde{A}_3$. Then $\tilde{A}_1, \tilde{A}_2, \tilde{A}_3$ is a partition of $[n]$ with $|\tilde{A}_1| = \frac{n}{3} + s_1 - s_2$, $|\tilde{A}_2| = \frac{n}{3}$, and $|\tilde{A}_3| = \frac{n}{3} - s_1 + s_2$.

If χ has a unique 3-collision $\{i, j, k\}$, then the 3-wise independence of f implies that

$$\begin{aligned} & \Pr_f \left(i \in \tilde{A}_1, j \in \tilde{A}_2, k \in \tilde{A}_3 \right) \\ &= \Pr_f \left(f(i) \leq \frac{n}{3} + s_1 - s_2, \frac{n}{3} + s_1 - s_2 < f(j) \leq \frac{2n}{3} + s_1 - s_2, f(k) > \frac{2n}{3} + s_1 - s_2 \right) \\ &\geq \left(\frac{1}{3} - o(1) \right)^3 = \frac{1}{27} - o(1). \end{aligned}$$

We assume that this holds when constructing the starting state. Otherwise, our construction fails and we try again. Let $\tilde{n}_2 := |\mathcal{P}(\tilde{A}_1, \tilde{A}_2)|$. We can assume that $\tilde{n}_2 \in \Omega(n)$ for the same reason we can always assume that χ has $\Omega(n)$ 2-collisions.

To begin, we create a uniform superposition $\left(\binom{\tilde{A}_1}{s_1}\right)^{-1/2} \sum_{I \in \binom{\tilde{A}_1}{s_1}} |Q(I)\rangle$ of sets of s_1 indices drawn from \tilde{A}_1 , stored in a skip-table, and query these indices. This uses s_1 queries and s_1 insertions, for a total time complexity of $\tilde{O}(s_1)$.

For $I \subset \tilde{A}_1$, let $H(I) := \{j \in \tilde{A}_2 : \exists i \in I, \chi_j = \chi_i\}$. Next, we search \tilde{A}_2 for indices in $H(I)$, assuming that $H(I)$ has size at least $\Omega(\frac{\tilde{n}_2 s_1}{n})$. The following lemma justifies this assumption.

Lemma 5.5. *Let I be a uniformly random subset of \tilde{A}_1 of size s_1 . Then $\Pr_I(|H(I)| \leq \frac{\tilde{n}_2}{n} s_1) < o(1)$.*

Proof. The random variable $|H(I)|$ has a hypergeometric distribution with mean $\mu = \frac{s_1 \tilde{n}_2}{2n/3 + s_1 - s_2} = \Theta(s_1)$. Using tail inequalities [Ska11, eq. 14] we have, for any constant $c \geq 1$,

$$\Pr(|H(I)| \leq \frac{1}{c} \mu) \leq \exp \left(-2 \left(\frac{\mu(1 - \frac{1}{c})}{s_1} \right)^2 s_1 \right) \leq e^{-\Theta(s_1)} = o(1). \quad \square$$

Thus we can then restrict our attention to the part of the state with $|H(I)| \geq \epsilon s_1$ for some constant $\epsilon \leq \frac{\tilde{n}_2}{n}$, as this part has $1 - o(1)$ of the weight. We can then perform the mapping

$$\left(\binom{\tilde{A}_1}{s_1}\right)^{-1/2} \sum_{I \in \binom{\tilde{A}_1}{s_1} : |H(I)| \geq \epsilon s_1} |Q(I)\rangle \mapsto \left(\binom{\tilde{A}_1}{s_1}\right)^{-1/2} \sum_{I \in \binom{\tilde{A}_1}{s_1} : |H(I)| \geq \epsilon s_1} |Q(I)\rangle \left(\binom{|H(I)|}{s_2}\right)^{-1/2} \sum_{J \in \binom{H(I)}{s_2}} |Q(J)\rangle$$

using s_2 applications of Grover search for a new element of $H(I)$. Each search requires $\tilde{O}(\sqrt{n/|H(I)|}) = \tilde{O}(\sqrt{n/s_1})$ iterations. As we find elements, we insert them into a skip-table, also separately recording the order in which we find the indices. To check if some i is in $H(I)$, but not already found, we

- look up i in the skip-table of indices already found;
- query χ_i ;
- compute $f(i)$; and
- look up χ_i in $Q(I)$.

Each of these operations has time complexity $\tilde{O}(1)$, for a total cost of $\tilde{O}(1)$ per iteration. The total cost of the s_2 rounds of search is $\tilde{O}(s_2 \sqrt{n/s_1})$. Finally, we must uncompute the order in which we found the indices of J . For each J , we have the state $|Q(J)\rangle 2^{-s_2/2} \sum_{\sigma \in \mathcal{S}_{s_2}} |\sigma(J)\rangle$, where \mathcal{S}_n is the symmetric group on n symbols. We can uncompute the order register in cost $\tilde{O}(s_2)$, completing the desired mapping.

Let I_1 be the elements of I for which we did not find a collision, and I_2 those elements of I for which we did find a collision. We can reversibly convert $|Q(I), Q(J)\rangle$ to $|Q(I_1), \{(i, j, \chi_i) : i \in I_2, j \in J, \chi_i = \chi_j\}\rangle$, where both sets are stored in a skip-table. We call the second set $Q(S_2)$. To accomplish this mapping, we do the following s_2 times, once for each $j \in J$:

- look up χ_j in $Q(I)$ to find $(i, \chi_i = \chi_j)$;
- insert (i, j, χ_i) into $Q(S_2)$; and

- delete (i, χ_i) from I and (j, χ_i) from J .

What remains in $Q(I)$ after performing these steps is exactly $Q(I_1)$. Each repetition costs $\tilde{O}(1)$, for a total cost of $\tilde{O}(s_2)$. Note that we can delete (i, χ_i) every time because all 2-collisions in $\tilde{A}_1 \times \tilde{A}_2$ are disjoint (by assumption), so $|I_1| = s_1 - s_2$. We also have $|H(I_1)| = |H(I)| - s_2$, again because all 2-collisions in $\tilde{A}_1 \times \tilde{A}_2$ are disjoint. Thus, after performing the full mapping, the part of the state under consideration is

$$\left(\binom{|\tilde{A}_1|}{s_1}\right)^{-1/2} \sum_{I_1 \in \binom{\tilde{A}_1}{s_1 - s_2}} \left(\binom{|H(I_1)| + s_2}{s_2}\right)^{-1/2} |Q(I_1)\rangle \sum_{S_2 \in \binom{\mathcal{P}(\tilde{A}_1 \setminus I_1, \tilde{A}_2)}{s_2}} |Q(S_2)\rangle.$$

Measuring the first register, containing some $|Q(I_1)\rangle$, gives the state

$$\left(\binom{|\mathcal{P}(\tilde{A}_1 \setminus I_1, \tilde{A}_2)|}{s_2}\right)^{-1/2} \sum_{S_2 \in \binom{\mathcal{P}(\tilde{A}_1 \setminus I_1, \tilde{A}_2)}{s_2}} |Q(S_2)\rangle$$

for some I_1 with probability at least $1 - o(1)$. Adding up the total cost, we find $S_C = \tilde{O}(s_1 + s_2\sqrt{n/s_1} + s_2) = \tilde{O}(s_1 + s_2\sqrt{n/s_1})$, since $n > s_1$.

This state is the correct starting state for the partition $A_1 = \tilde{A}_1 \setminus I_1$, $A_2 = \tilde{A}_2$, $A_3 = \tilde{A}_3 \cup I_1$, which is clearly a tripartition with $|A_1| = |A_2| = |A_3| = n/3$. Furthermore, for a 3-collision $\{i, j, k\}$, assuming $i \in \tilde{A}_1, j \in \tilde{A}_2, k \in \tilde{A}_3$ (which happens with constant probability), the only way we can fail to have $i \in A_1, j \in A_2, k \in A_3$ is if $i \in I_1$. Although I_1 is not uniformly distributed, $\Pr(I_1 | H(I_1) = h)$ is uniform for any h . Furthermore, $\Pr(i \in I_1 | H(I_1) = h) = o(1)$ for any fixed h , since $|I_1| \ll n$. Thus, we have $\Pr_{A_1, A_2, A_3}(i \in A_1, j \in A_2, k \in A_3) = \Omega(1)$.

Finally, to store the tripartition A_1, A_2, A_3 , we need to keep f , as well as I_1 , which we store in a skip-table. This takes space $\tilde{O}(1) + \tilde{O}(|I_1|) = \tilde{O}(s_1)$. To compute which of A_1, A_2 , or A_3 contains an index i , we first compute $f(i)$, and then (possibly) look up i in I_1 , each of which costs $\tilde{O}(1)$. \square

Theorem 5.6 (Costs of the update walk $S', \frac{1}{\sqrt{\varepsilon'}}(\frac{1}{\sqrt{\delta'}}U' + C')$). *The update walk has time complexities $S' = \tilde{O}(s_1)$ and $\frac{1}{\sqrt{\varepsilon'}}(\frac{1}{\sqrt{\delta'}}U' + C') = \tilde{O}(\sqrt{nm/s_1})$.*

Proof. This follows from Theorem 4.3 and our encoding of a triple $(Q(S_2), Q(S_1), Q(S'_1))$. The implementation is nearly identical to the time-efficient Element Distinctness algorithm of [Amb04], except that we store an extra skip-table containing the set $Q(\mathcal{P}(S_1))$. However, insertion and deletion may still be performed in poly-logarithmic time. To insert i into S_1 , we must look up χ_i in $Q(S_1)$ to see if we have a new collision in $Q(S_1)$, i.e., if there is some (j, χ_i) already in $Q(S_1)$ such that $(i, j) \in (A_1 \times A_2) \cup (A_2 \times A_1)$. If there is such a j , then we insert (i, j, χ_i) into $Q(\mathcal{P}(S_1))$ if $(i, j) \in A_1 \times A_2$ or (j, i, χ_i) into $Q(\mathcal{P}(S_1))$ else. Finally, we insert (i, χ_i) into $Q(S_1)$. This involves a constant number of skip-table insertions and lookups, so its cost is still poly-logarithmic. We can delete some $i \in S_1$ by running this operation in reverse. Thus, the update and setup cost are clearly the same as [Amb04]. From the proof of Theorem 4.3, we have

$$S' = \tilde{O}(s_1) \quad U' = \tilde{O}(1) \quad \delta' = \Omega(1/s_1) \quad \varepsilon' = \Omega(1).$$

To check if S_1 is marked, we simply read the counter storing the size of $Q(\mathcal{P}(S_1))$ and check if it is at least m , in time $C' = O(1)$. Thus, we have

$$S' = \tilde{O}(s_1) \quad \frac{1}{\sqrt{\varepsilon'}} \left(\frac{1}{\sqrt{\delta'}} U' + C' \right) = \tilde{O} \left(\sqrt{\frac{nm}{s_1}} \right). \quad \square$$

Since $|\psi(S_2, S'_2)\rangle = |\psi(S'_2, S_2)\rangle$ for all edges (S_2, S'_2) by Lemma 4.4, we have the following.

Theorem 5.7 (GARBAGE SWAP cost). *We can implement the GARBAGE SWAP in time $\tilde{O}(m)$.*

Proof. The GARBAGE SWAP is the operation that acts, for any edge (S_2, S'_2) , as

$$|Q(S_2), Q(S'_2)\rangle |\psi(S_2, S'_2)\rangle \mapsto |Q(S'_2), Q(S_2)\rangle |\psi(S'_2, S_2)\rangle.$$

By Lemma 4.4, we need only consider the cost of $|Q(S_2), Q(S'_2)\rangle \mapsto |Q(S'_2), Q(S_2)\rangle$. Recall that $|Q(S_2), Q(S'_2)\rangle$ is encoded as $|Q(S_2)\rangle |Q(S_2 \setminus S'_2)\rangle |Q(S'_2 \setminus S_2)\rangle$, with each of the three parts encoded as a skip-table. Since (S_2, S'_2) is an edge, we have $|S_2 \setminus S'_2| = |S'_2 \setminus S_2| = m$. Thus, we can perform the mapping to $|Q(S'_2)\rangle |Q(S'_2 \setminus S_2)\rangle |Q(S_2 \setminus S'_2)\rangle$ by performing m insertions and m deletions on $Q(S_2)$ to get $Q(S'_2)$. \square

Theorem 5.8 (LOCAL DIFFUSION WITH GARBAGE cost). *We can implement the LOCAL DIFFUSION WITH GARBAGE with time complexity $\tilde{O}(m)$.*

Proof. We consider each of the three steps from the proof of Theorem 4.6. In step 1, we create a superposition over sets of m values in $Q(S_2)$ using m superposition accesses to the skip-table storing $Q(S_2)$. In step 2 we create a superposition over sets of m values in $Q(\mathcal{P}(S_1))$ using m superposition accesses to the skip-table storing $Q(\mathcal{P}(S_1))$. Finally, in step 3 we perform m lookups in the skip-table storing $Q(J) = Q(S'_2 \setminus S_2)$ and $2m$ deletions from the skip-table storing $Q(S_1)$. The total cost of this is $\tilde{O}(m)$.

It is also clear from the proof of Theorem 4.6 that we move from a superposition of correctly encoded triples $(Q(S_2), Q(S_1), 0)$ (where the 0 corresponds to the coin of S_1) to a superposition of correctly encoded triples $(Q(S_2), Q(S'_2), Q(\tilde{S}_1))$. \square

Corollary 5.9 (LOCAL DIFFUSION and DATABASE SWAP cost T). *The family $(P^{S_2}, M^{S_2}, d^{S_2})_{S_2 \in X}$ implements the LOCAL DIFFUSION and DATABASE SWAP of (P, Q) with time complexity $\mathsf{T} = \tilde{O}(m)$.*

Proof. This is immediate from Theorems 5.7 and 5.8. \square

Theorem 5.10 (Checking cost C). *We can implement the checking reflection in time $\mathsf{C} = \tilde{O}(\sqrt{n})$.*

Proof. To check if a vertex S_2 is marked, we search for an index $k \in A_3$ such that there exists $(i, j) \in S_2$ such that $\{i, j, k\}$ is a 3-collision. Each time we check if a particular k has this property, we query k and look up χ_k in $Q(S_2)$ in time $\tilde{O}(1)$. \square

Theorem 5.1 now follows. All costs are the same as their query complexities from Section 4, with the exception of $\mathsf{T} = \tilde{O}(m) = \tilde{O}(\frac{s_2^2}{n})$, but this does not change the asymptotic complexity. Plugging the values from Theorems 5.4, 5.6, and 5.10 and Corollary 5.9 into the nested update cost expression from Theorem 3.3, we have

$$\begin{aligned} \mathsf{S}_C + \mathsf{S}' + \frac{1}{\sqrt{\varepsilon}} & \left(\frac{1}{\sqrt{\delta}} \left(\frac{1}{\sqrt{\varepsilon'}} \left(\frac{1}{\sqrt{\delta'}} \mathsf{U}' + \mathsf{C}' \right) + \mathsf{T} \right) + \mathsf{C} \right) \\ &= s_1 + s_2 \sqrt{\frac{n}{s_1}} + s_1 + \sqrt{\frac{n_2}{s_2}} \left(\sqrt{\frac{s_2}{m}} \left(\sqrt{\frac{nm}{s_1}} + m \right) + \sqrt{n} \right) \\ &= s_1 + s_2 \sqrt{\frac{n}{s_1}} + \frac{n}{\sqrt{s_1}} + \sqrt{n_2 \frac{s_1^2}{n}} + \frac{n}{\sqrt{s_2}} = \tilde{O} \left(s_1 + s_2 \sqrt{\frac{n}{s_1}} + \frac{n}{\sqrt{s_1}} + \frac{n}{\sqrt{s_2}} \right), \end{aligned}$$

which is still optimized by setting $s_1 = n^{5/7}$ and $s_2 = n^{4/7}$, giving time complexity $\tilde{O}(n^{5/7})$.

6 Conclusion and Future Directions

We have shown that the quantum walk search framework of [MNRS11] can be extended to allow a data function that depends on both the vertex and the coin, provided certain costs are accounted for. This extension allows us to implement nested updates, although there may be other applications of this new framework, as it more generally allows us to consider updates with garbage resulting from any type of update subroutine. Nested updates provide another tool for quantum walk algorithms analogous to the nested checking of [JKM13], and we hope that these tools will facilitate further upper bounds on both time and query complexity.

It remains an open problem to improve the $\tilde{O}(n^{\frac{k}{k+1}})$ time complexity upper bound for k -Distinctness for $k > 3$. The 3-Distinctness upper bound of [Bel13] can be extended to a general k -Distinctness upper bound by coming up with an efficient procedure for constructing a starting state that generalizes our $|\pi\rangle_D^0$ [Bel13]. Efficiently constructing this generalized starting state would be a necessary, but not sufficient, condition for generalizing our upper bound to $k > 3$.

7 Acknowledgments

We thank Aleksandrs Belovs for helpful discussions about our two different approaches to this problem.

References

- [Amb04] A. Ambainis. Quantum walk algorithm for element distinctness. In *Proceedings of 45th IEEE Symposium on Foundations of Computer Science*, pages 22–31, 2004.
- [AS04] S. Aaronson and Y. Shi. Quantum lower bounds for the collision and element distinctness problems. *Journal of the ACM*, 51:595–605, 2004.
- [BDH⁺05] H. Buhrman, C. Dürr, M. Heiligman, P. Høyer, M. Santha, F. Magniez, and R. de Wolf. Quantum algorithms for Element Distinctness. *SIAM Journal on Computing*, 34(6):1324–1330, 2005.
- [Bel12a] A. Belovs. Learning-graph-based quantum algorithm for k -distinctness. In *Proceedings of 53rd IEEE Symposium on Foundations of Computer Science*, pages 207–216, 2012.
- [Bel12b] A. Belovs. Span programs for functions with constant-sized 1-certificates. In *Proceedings of 44th Symposium on Theory of Computing*, pages 77–84, 2012.
- [Bel13] A. Belovs. Quantum walks and electric networks, 2013. [arXiv:1302.3143](https://arxiv.org/abs/1302.3143).
- [CK11] A. M. Childs and R. Kothari. Quantum query complexity of minor-closed graph properties. In *Proceedings of 28th Symposium on Theoretical Aspects of Computer Science*, pages 661–672, 2011.
- [JKM13] S. Jeffery, R. Kothari, and F. Magniez. Nested quantum walks with quantum data structures. In *Proceedings of 24th ACM-SIAM Symposium on Discrete Algorithms*, pages 1474–1485, 2013.

- [LMR⁺11] T. Lee, R. Mittal, B. Reichardt, R. Spalek, and M. Szegedy. Quantum query complexity of state conversion. In *Proceedings of 52nd IEEE Symposium on Foundations of Computer Science*, pages 344–353, 2011.
- [LMS11] T. Lee, F. Magniez, and M. Santha. A learning graph based quantum query algorithm for finding constant-size subgraphs, 2011. [arXiv:1109.5135](#).
- [LMS13] T. Lee, F. Magniez, and M. Santha. Improved quantum query algorithms for triangle finding and associativity testing. In *24th ACM-SIAM SODA*, pages 1486–1502, 2013.
- [MNRS11] F. Magniez, A. Nayak, J. Roland, and M. Santha. Search via quantum walk. *SIAM Journal on Computing*, 40(1):142–164, 2011.
- [MNRS12] F. Magniez, A. Nayak, P. Richter, and M. Santha. On the hitting times of quantum versus random walks. *Algorithmica*, 63(1):91–116, 2012.
- [Rei11] B. Reichardt. Reflections for quantum query algorithms. In *Proceedings of 22nd ACM-SIAM Symposium on Discrete Algorithms*, pages 560–569, 2011.
- [Ska11] M. Skala. Hypergeometric tail inequalities: ending the insanity, 2011.
- [Sze04] M. Szegedy. Quantum speed-up of Markov chain based algorithms. In *Proceedings of 45th IEEE Symposium on Foundations of Computer Science*, pages 32–41, 2004.
- [WC81] M. N. Wegman and L. Carter. New hash functions and their use in authentication and set equality. *Journal of Computer and System Sciences*, 22(3):265–279, 1981.